MASSIVELY PARALLEL VISUALIZATION ON LEADERSHIP COMPUTING RESOURCES

Tom Peterka[1], Robert Ross[1], Hongfeng Yu[2], and Kwan-Liu Ma[3]

[1]Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Ave., Argonne, IL 60439-4803; [2]Sandia National Laboratories, California, PO Box 969, Livermore, CA 94551-0969; [3] Department of Computer Science, University of California at Davis, One Shields Ave., Davis, CA 95616-8562

INTRODUCTION

Supercomputer size, computing power, and computational results are rapidly approaching petascale proportions. For example, with the recent deployment of the Blue Gene/P (BG/P) at Argonne's Leadership Computing Facility [1], scientists can compute simulations at a rate of hundreds of teraflops (trillions of floating point operations per second) and produce datasets hundreds of terabytes in size. This will soon grow to petaflops and petabytes. Performing post-simulation analysis and visualization operations using a dedicated analysis computer cluster, as we do now, will not scale with the rapidly expanding computational power of the supercomputers performing the simulations. Studying how simulation, analysis, and visualization can be collocated on the same leadership system has the potential benefits of eliminating data transfers and lays the groundwork for performing *in situ* visualization – viewing results of a simulation as they unfold [2]. In this research, we examine how a popular scientific visualization operation, volume rendering, can be massively parallelized on BG/P and evaluate its performance when scaling from 64 to 32,768 processor cores.

METHOD

Volume rendering is the process of classifying data values according to color and opacity and then casting rays from the viewpoint through each image pixel and the dataset, in our case a time-varying 3D scalar field from a supernova simulation. Colors and opacities of each pixel are accumulated as these rays travel through the data volume. To parallelize this algorithm, we subdivide the dataset into regions and assign each region to a different processor. Each processor reads its portion of the dataset from storage and carries out the above ray casting algorithm. Then, these partial results are composited together via a mass-exchange of completed subimages to form a final completed image. We examine the performance of the algorithm at large system scales, identifying and solving bottlenecks along the way.

RESULTS

Figure 1 shows that large visualizations are I/O bound. For example, at 512 processors, 75% of the total time is spent reading the time-step while 25% of the time is spent rendering it. This disparity grows with increasing numbers of processors. Compositing time is minimal until 4096 processors.

These visualization experiments are the largest system scale ever used for this type of application. At tens of thousands of processors, the distribution of I/O, rendering, and compositing time argues for a leadership-class system for visualization; the storage and communication subsystems are high-bandwidth and low-latency, capable of serving the needs of the massive numbers of computational nodes. As part of this research, we have increased image quality, extended the algorithm to generate stereoscopic images, improved I/O performance, mitigated the remaining effects of I/O, and extended the range of effective composting by an order of magnitude [3, 4]. We conclude from numerous tests that leadership systems such as BG/P can have potential benefits not only for computation but for analysis and visualization as well. This novel approach to visualization can eliminate terabytes of data movement between systems, and sets the stage for *in situ* visualization, where these tasks occur during the same time as a running simulation.
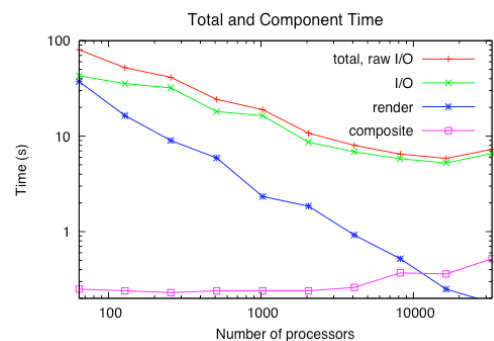

Figure1: Individual component and total timing results across a range of processor cores.

ACKNOWLEDGMENTS

REFERENCES

[1] Argonne Leadership Computing Facility. http://www.alcf.anl.gov/ 2008.

[2] MA, K.-L., WANG, C., YU, H. TIKHONOVA, A.: In-Situ Processing and Visualization for Ultrascale Simulations. *Journal of Physics*, 78, (June 2007).

[3] PETERKA, T., ROSS, R., YU, H., MA, K.-L. KOOIMA, R.: Coupling Massively Parallel Volume Rendering with Autostereoscopic 3D Display Environments. *Proceedings of International Workshop on Super Visualization (IWSV'08)*, Kos, Greece, (June 2008).

[4] PETERKA, T., YU, H., ROSS, R. MA, K.-L.: Parallel Volume Rendering on the IBM Blue Gene/P. *Proceedings of Eurographics Parallel Graphics and Visualization Symposium 2008*, Crete, Greece, (April 2008).